arXiv:1612.02522v1 [cs.NE] 8 Dec 2016

Geometric Decomposition of Feed Forward Neural Networks

Sven Cattell

December 9, 2016

Abstract

There have been several attempts to mathematically understand neural networks and many more from biological and computational perspectives. The field has exploded in the last decade, yet neural networks are still treated much like a black box. In this work we describe a structure that is inherent to a feed forward neural network. This will provide a framework for future work on neural networks to improve training algorithms, compute the homology of the network, and other applications. Our approach takes a more geometric point of view and is unlike other attempts to mathematically understand neural networks that rely on a functional perspective.

1 Introduction

In the last decade deep learning has exploded. With many applications to image recognition, natural language processing, and many other pattern recognition tasks that have traditionally had poor performance with other machine learning techniques seem to be finally cracked. However, deep learning is rather poorly understood mathematically. In the original papers the intuition provided is compelling and has proved to be useful in developing the technique [6][1][8]. However, this needs some mathematical backing in order to properly use and improve the techniques.

This paper aims to give a description of the parts of a neural network and their functions. In order to simplify the structure of a fully connected feed forward neural network to one that is easily understood we will mainly work with the step edge activation function. This may seem like a step back from the modern use of smooth or piecewise linear activation functions, but if the activation function we want to train against is a approximation of the step edge function (like the sigmoidal function), then we can still make strong inferences of the general structure. This paper focuses on the geometry of the network so, unless otherwise noted, we are using the step-edge activation function everywhere. The main theorem of this paper is:

Theorem 1.1. *A binary classification neural network is an indicator function on a union of some regions in* $\mathbb{R}^n - \bigcup_{P\in\mathcal{A}} P$*, where* $\mathcal{A}$ *is defined by the first layer.*

This is not similar to previous work in mathematically understanding shallow neural networks of Kolmogorov [7], Funahashi [5]or Sprecher [12]. They prove that there is some network that can approximate a given function. However this is non-constructive and does not apply to deep networks. The more recent results of [11] and [2] construct well understood networks of arbitrary precision. This paper is intended to give a more general structure theory that can be applied to build networks or manipulate them.

2 Background

2.1 Neural Networks

This part of the background is intended for mathematicians unfamiliar with neural networks. Feed forward neural networks are easily expressed as a composition of linear functions, with a non-linear *activation function*. We will denote the activation function by $\mathfrak{h}$. It is always defined as a single variable function. However, it is regularly referred to as being defined on $\mathbb{R}^n$, not just $\mathbb{R}$, this is just $\mathfrak{h}^{\oplus n}$, or $\mathfrak{h}$ on each coordinate of $\mathbb{R}^n$.

Definition 1. *A* classification neural network *of height* h *with* n *inputs and* k *outputs is a function* $N : \mathbb{R}^n \to \mathbb{R}^k$ *such that it can be written as a composition of* $h+1$ *functions of the form* $\mathfrak{h}(A(x)+b)$ *where* A *is some matrix* b *is an offset*

vector called layers. *In other words there exists* $h+1$ *matrices* A_i*, with offset vectors* b_i*, and:*

$$N(x) = \mathfrak{h}(A_{h+1}\mathfrak{h}(A_h(\dots\mathfrak{h}(A_2\mathfrak{h}(A_1(x)+b_1)+b_2)\dots)+b_h)+b_{h+1})$$

Each A_i *is a matrix of size* $k_i \times k_{i-1}$ *and* b_i *is a* k_i*-tuple, we say the layer* i *has* k_i *nodes.*

Of course, this is a rather obtuse definition, it is usually expressed as a directed graph. The nodes are organized into layers with directed edges pointing up the layers. These edges are given weights and each node, other than the input nodes, are a linear sum of all nodes that have an edge pointing towards the aforementioned node. Sometimes the offset is represented by an additional input node who is connected to all the nodes in higher layers and has a fixed value of 1, but here we shall think of the offset as a part of each node that isn't a input node.

For simplicity, we will mainly discuss a neural network with one output. All the results generalize quite easily to multiple outputs, so they are mostly omitted.

2.2 Hyperplane Arrangements

This section of the background is intended for those familiar with neural networks, but unfamiliar with the language of hyperplane arrangements. It is intended to be a short summary of the necessary language required for the main result. The following is mostly lifted from [10], which has much more detail that we may provide here.

An *arrangement of hyperplanes*, $\mathcal{A}$ in $\mathbb{R}^n$ is a finite set of codimension 1 affine subspaces. The *regions* between the hyperplanes refer to the components of $\mathbb{R}^n - \bigcup_{P\in\mathcal{A}} P$. An arrangement is in *general position* if a small perturbation of the hyperplanes in the arrangement does not change the number of hyperplanes. For example two parallel lines in $\mathbb{R}^2$ are not in general position. An arbitrarily small rotation of one of the lines will result in the number of regions increasing from 3 to 4. The other essential structure inherent in hyperplane arrangements

Definition 2. *The intersection poset of an arrangement of* k *hyperplanes,* $\mathcal{A}$ *is the set*

$$L(\mathcal{A}) = \{x \subset \{1,\dots,k\} | \mathbb{R}^n \cap \bigcap_{i\in x} P_i \neq \emptyset\}$$

where $P_i \in \mathcal{A}$*. This is equipped with the partial order* $x \leq y$ *if* $x \subset y$*.*

A poset is a set equipped with a partial relation, i.e. if $x, y \in X$ then we may have $x \neq y$, $x \not< y$ and $y \not< x$. In a poset, an element y *covers* x, or $x \lessdot y$, if there is no z different from x and y such that $x < z < y$. We can draw the Hasse diagram of a poset by drawing a node for each element of the poset and a directed edge from the node for x to the node for y if $x \lessdot y$.

3 Characterization of a Perceptron

If we examine the first layer of a neural network each node corresponds to a weighted sum of the input nodes with an offset. We can represent the weighted sum by a vector v, and the offset by b. Then the node is "active" on the point x, and outputting a 1 if $v \cdot x > b$, and 0 otherwise. This is the same as the indicator function on the positive component, U, of $\mathbb{R}^n - P$ where P is the associated hyperplane the following set:

$$U = \{x \in \mathbb{R}^n | v \cdot x > b\}.$$

So, another way of expressing a node on the first layer is by the indicator function on the set U, $\mathbb{1}_U$.

Therefore the input, under the step edge activation function, to the second layer of the network is the output of a collection of n indicator functions. The second layer has a binary input on n variables for each point in the underlying space. Therefore all it can do is make decisions based on which side of each hyperplane the given input point is. We will use $\mathcal{A}$ to denote a collection of hyperplanes in R^n. This divides $\mathbb{R}^n$ into regions, $\mathcal{R}$, the connected components of $\mathbb{R}^n - \bigcup_{P \in \mathcal{A}} P$. The second layer is only aware of which region you are in as each region produces a unique signature output of the first layer. The subsequent layers are a means of making a choice of which regions to include, each layer past the first will amount a process we call a weighted union of the sets associated to the nodes of the previous layer.

3.1 Regions in a Polarized Arrangement

A plane in $\mathbb{R}^n$ is defined by a normal vector v and a offset value b. If we require that the normal vector be of length 1, then there are two normal vectors to choose from. Both work, but they define a different orientation of the hyperplane. Therefore, rather than just having two components of $\mathbb{R}^n$ we can now distinguish between them.

Definition 3. *A polarization of a hyperplane in $\mathbb{R}^n$ for the plane P with normal vector v and offset b are the two sets:*

$$\mathcal{R}_P^+ = \{x | x \cdot v + b > 0\}$$

$$\mathcal{R}_P^- = \{x | x \cdot v + b < 0\}$$

We call $\mathcal{R}_P^+$ the positive side and U_P^- the negative side of the plane.

We usually index our hyperplanes by the numbers $\{1, 2, \ldots k\}$ so we write $\mathcal{R}_i^+$ for the positive side of the ith hyperplane in this case. Nodes in the first layer of the neural network are the *hyperplane layer* as they determine a polarized arrangement of hyperplanes.

For a polarizes arrangement of hyperplanes $\mathcal{A}$ indexed by a set $I = \{1, 2, \ldots k\}$ we can define and label the regions.

Definition 4. *The regions, $\mathcal{R}$ of a polarization arrangement of hyperplanes $\mathcal{A}$ is the set of convex polytopes formed by taking all possible intersections of the positive and negatives sides of an ordered set of k plane partitions. Each region is labeled by a $J \subset I$ where*

$$\mathcal{R}_J = \bigcap_{j \in J} \mathcal{R}_j^+ \cap \bigcap_{j \in I-J} \mathcal{R}_j^-$$

We are using a similar definition to [9], but with a subset of $\{1, \ldots, n\}$ instead of a ordered n-tuple. The two methods are isomorphic, but we choose to use sets for now. There are 2^k different labelings one could have for a hyperplane arrangement, depending on the polarization. See figure 3.1 for an example.

Note, not all possible labellings are non-empty regions as the largest possible number of regions in for an arrangement of k hyperplanes $\mathbb{R}^n$ is:

$$r(\mathcal{A}) = 1 + k + \binom{k}{2} + \ldots \binom{k}{n}.$$

Which is 2^n if $k = n$, but when $k > n$ is strictly less than 2^k. This is when the regions are in general position, that is we can change the values of all the normal vectors and offsets by a small amount and the number of regions will not change. The initial layer of a perceptron will most likely be in general position. We call an index $J \subset I$ trivial if $\mathcal{R}_J = \emptyset$.

3.2 Weighted Unions and Selection Layers

Layers of the neural network after the first layer will amount to a weighted union of the sets associated to the previous layer. The second layer is a weighted union of the positive sides of the polarized hyperplanes generated by the first layer. This results in each node being equivalent to a union of the regions of the arrangement. Thus we call a layer a *selection layer* if it is not the first layer. We start with a set level definition of what each node in a selection layer:

Definition 5. *A weighted union of subsets of X, $\{U_i\}_{i \in I}$, with weights a_i and offset b is*

$$\bigcup_{i \in I} (a_i U_i) - b = \left\{ x | \mathfrak{h}(\sum_{i \in I} a_i \mathbb{1}_{U_i}(x) - b) > 0 \right\}.$$

The characteristic function is defined on X.

It is clear that this is just taking the output of a node and converting it back to the associated set. In order to manipulate a selection node we need a clear understanding of the weighted union. For a single set $U \subset X$, we can take the compliment of U by the weighted union: $-1U - \frac{1}{2}$. We can also take the union and intersection of two or more sets:

$$\bigcap_{i=1}^{n} U_i = \bigcup_{i=1}^{n} 1U_i - (n - \frac{1}{2}), \qquad \bigcup_{i=1}^{n} U_i = \bigcup_{i=1}^{n} 1U_i - (\frac{1}{2}).$$

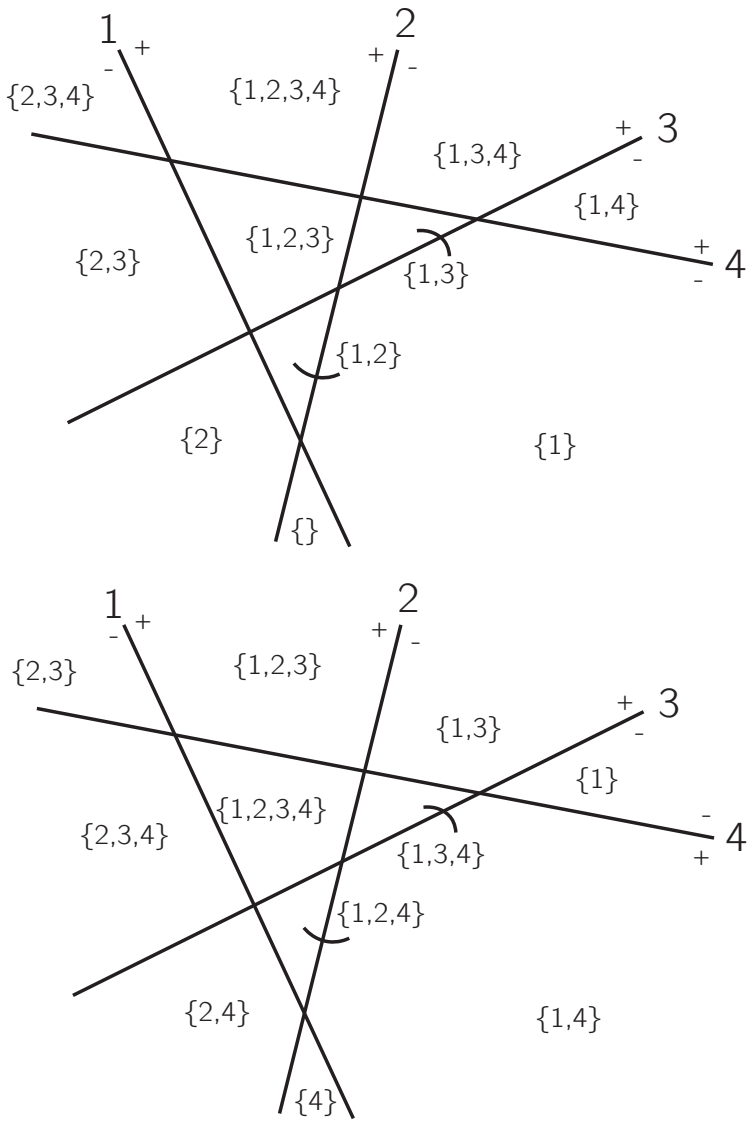

Figure 1: Two different labellings of the same regions of an arrangement of hyperplanes. The difference is the plane labeled by 4 changes polarity. Note, no one region needs to be labeled with $\emptyset$ unlike the labeling from [3]

The following lemma demonstrates some of the limitations of the weighted union:

Lemma 3.1. *A weighted union of sets U_i can be written as a union a subset of all intersections of those sets and their complements.*

Proof. Let $W = \bigcup_{i\in I}(a_i U_i) - b$ be a weighted union. For each a_i,

$$a_i \mathbb{1}_{U_i}(x) = -a_i \mathbb{1}_{X-U_i}(x) + 2$$

therefore without loss of generality we may assume for all i, $a_i > 0$, else replace a_i with $-a_i$, add a_i to b and replace U_i with $X - U_i$. Let $\tilde{b}$ be the adjusted offset For $J \subset I$ Let

$$U_J = \bigcap_{j\in J} U_j \cap \bigcap_{j\in I-J} (X - U_j),$$

The map

$$\sigma : P(I) \to \mathbb{R}, \quad \sigma(J) \mapsto \sum_{j\in J} a_j,$$

is order preserving as all a_i are non-negative. Let $\{x > \tilde{b}\}$ be the obvious shorthand for $\{x \in \mathbb{R} | x > \tilde{b}\}$, it is easy to show:

$$W = \bigcup_{J\in\sigma^{-1}(\{x>\tilde{b}\})} U_J.$$

□

Definition 6. *For a collection of sets $\{U_i\}_{i\in I}$, and for each $J \in P(I)$ the we get a* region *of the*

$$U_J = \bigcap_{j\in J} U_j \cap \bigcap_{j\in I-J} (X - U_j)$$

The indexing of the region U_J by J is the standard indexing

We can see from the proof that if we want to find out what set operations are possible with the weighted union we can restrict ourselves to a union of some intersections of our sets. This can be seen through the light of a n-ary logical operation that is strictly composed out of a series of 'or's on a collection of 'and's. This cannot produce all n-ary boolean statements [4].

The proof provides a way to translate the weights and offset to a map on the poset $P(I)$. However the mapping is very dependent on the signs of the original weights. Given two different weighted unions onto the same space we not only have a σ_1 and a σ_2, but we also have two different sign corrections. This means that it's more difficult to compare two selection nodes using this technique. We can however remove the sign correction for the polarization after we have determined the which intersections were selected.

To standardize the polarization we define the following. For any map $p : I \to \{+, -\}$ we may define a self map $r_p : P(I) \to P(I)$ where $j \in r_p(J)$ if

$j \in J$ and $p(j) = +$ or if $j \notin J$ and $p(j) = -$. r_p is clearly a bijection for all p. For a collection of sets $\{U_i\}_{i\in I}$ and each p we get a different polarization of all possible intersections of the sets and their complements. The p polarization of the regions of $\{U_i\}_{i\in I}$ is equivalent to the standard polarization over $\{U_i^{p(i)}\}_{i\in I}$ where $\mathcal{U}_i^+ = U_i$ and $U_i^- = X - U_i$.

For the weights $\{a_i\}_{i\in I}$ and offset b the polarization of the regions is dependent on p such that $p(i) = +$ if $a_i \geq 0$ and $p(i) = -$ if $a_i < 0$. Therefore the polarization of the regions in a weighted union, $\sigma(\{x > \tilde{b}\})$, is over the p polarization of all regions. To convert our polarization to be over the standard polarization we may take the image, $r_p(\sigma(\{x > \tilde{b}\})$.

Definition 7. *A* selection $\mathcal{S}$ *of regions of a collection of sets indexed by* $I_{\mathcal{S}} \subset P(I)$ *to be the union of the regions:*

$$\mathcal{S} = \bigcup_{J\in\sigma^{-1}(\{x>\tilde{b}\})} U_J.$$

3.3 Characterization of a Perceptron

We now will characterize a perceptron as a series of weighted unions on top of a polarized arrangement of hyper planes. Figure 3.3 shows the various subnetworks in a neural network with 2 inputs, 4 planar nodes in the first layer, 2 selection nodes in the second layer and a selection node in the output layer.

Lemma 3.2. *A union or intersection of two selections of a plane partition set is another selection of that plane partition set.*

Theorem 3.3. *A binary classification neural network is an indicator function on a union of some regions of an arrangement of hyperplanes* $\mathcal{A}$*, where* $\mathcal{A}$ *is defined by the first layer.*

Proof. We will induct on the hidden height of a neural network. For the inductive step let N be a neural network of height $k + 1$. Each node in the final hidden layer of the neural network is a neural network of height k, let U_i be the selection of regions for the ith node. Let the final node have weight a_i from the ith node with offset b_{k+1}:

$$N(x) = \mathfrak{h}\left(\sum_i a_i \mathbb{1}_{U_i}(x) - b_{k+1}\right) = \mathbb{1}_{\bigcup a_i U_i - b_{k+1}}(x).$$

As the weighted union is a finite number of intersections and unions on its the inputs and the inputs are all selections of the same regions by lemma 3.2 the weighted union is. □

We can see that if every selection in the penultimate layer pairs two regions i.e. they are both selected for or selected against together then the final selection node cannot separate the two.

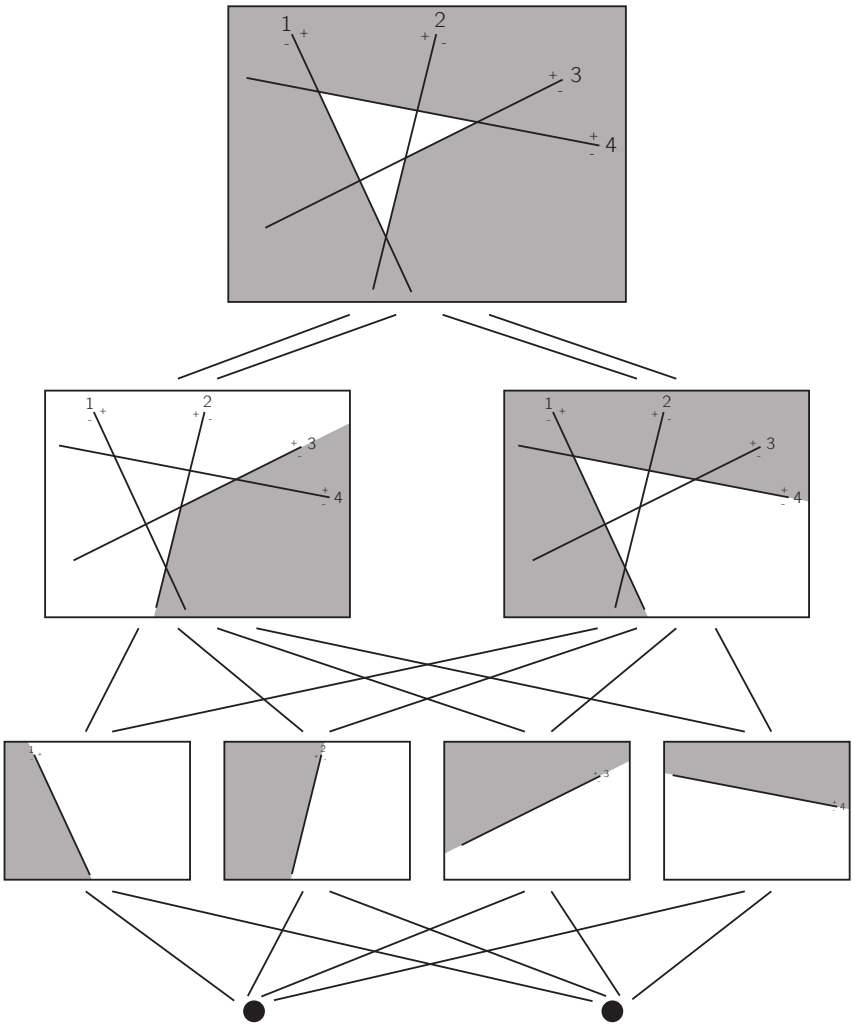

Figure 2: A neural network with 4 plane nodes, 2 selection nodes and an output, which is also a selection node. For each node we have highlighted the area where the associated neural subnetwork is positive in grey.

We can see that a neural network with k outputs is going to be k selections on the regions for the arrangement of hyperplanes. However, they are not independent from each other. The selections made by the last hidden layer affect the final possible selection.

References

[1] Yoshua Bengio, Nicolas L Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. In *Advances in neural information processing systems*, pages 123–130, 2005.

[2] C. K. Chui and H. N. Mhaskar. Deep nets for local manifold learning. *ArXiv e-prints*, July 2016.

[3] Paul H. Edelman. A partial order on the regions of r^n dissected by hyperplanes. *Transactions of the American Mathematical Society*, 283(2):617–631, 1984.

[4] H.B. Enderton. *A Mathematical Introduction to Logic*. Harcourt/Academic Press, 2001.

[5] K. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Netw.*, 2(3):183–192, May 1989.

[6] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

[7] A. N. Kolmogorov. On the representation of continuous functions of several variables as superpositions of continuous functions of one variable and addition. *Proceedings of the USSR Academy of Sciences*, 114:369–373, 1957.

[8] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[9] S. Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106(2):286 – 303, 1993.

[10] Ezra Miller. *Geometric combinatorics*. American Mathematical Society Institute for Advanced Study, Providence, R.I. Princeton, N.J, 2007.

[11] U. Shaham, A. Cloninger, and R. R. Coifman. Provable approximation properties for deep neural networks. *ArXiv e-prints*, September 2015.

[12] David A Sprecher. On the structure of continuous functions of several variables. *Transactions of the American Mathematical Society*, 115:340–355, 1965.